

What is claimed is:

- 1 1. A method of determining lower and upper bounds for a minimum cost
2 comprising the steps of:
 - 3 solving an integer program using a relaxation of binary variables to
4 determine the lower bound, the binary variables having values between
5 zero and one comprising a first subset;
 - 6 for the binary variables in the first subset and until no binary variables
7 remain in the first subset, iteratively performing the steps of:
 - 8 rounding up a first binary variable having a lowest ratio of a
9 cost penalty to a performance reward; and
 - 10 until no binary variables remain in a second subset, iteratively
11 performing the steps of:
 - 12 determining the binary variables in the first subset that
13 may be rounded down without violating a performance
14 constraint, thereby forming the second subset;
 - 15 rounding down one or more second binary variables in
16 the second subset having a zero performance reward; and
 - 17 rounding down a third binary variable in the second
18 subset having a highest ratio of a cost reward to the
19 performance reward if none of the binary variables in the
20 second subset have the zero performance reward; and
 - 21 determining the upper bound according to the binary variables having
22 binary values.
 - 1 2. The method of claim 1 wherein the integer program comprises the
2 performance constraint and an objective of minimizing a cost.
 - 1 3. The method of claim 1 wherein the integer program models a data placement
2 problem.
 - 1 4. The method of claim 3 wherein the data placement problem seeks to minimize
2 the cost of placing data objects onto nodes of a distributed storage system while
3 meeting a performance requirement for a workload.

- 1 5. The method of claim 1 wherein the step of rounding up the first binary
- 2 variable within the first subset further comprises calculating the cost penalty and
- 3 the performance reward.
- 1 6. The method of claim 5 wherein the step of rounding down the one or more
- 2 second binary variables within the second subset further comprises calculating the
- 3 performance reward.
- 1 7. The method of claim 6 wherein the step of rounding down the third binary
- 2 variable within the second subset further comprises calculating the cost reward.
- 1 8. A method of determining bounds for a minimum cost comprising the steps of:
2 solving an integer program using a relaxation of binary variables to
3 determine a lower bound for the minimum cost, the relaxation allowing the
4 binary variables to take values over the range of zero to one, a first subset
5 of the binary variables comprising the binary variables having values
6 between zero and one, the integer program modeling a data placement
7 problem which seeks to minimize a cost of placing data objects onto nodes
8 of a distributed storage system while meeting a performance requirement
9 for a workload;
10 until no binary variables remain in the first subset, iteratively
11 performing the steps of:
12 calculating a cost penalty and a performance reward for each of
13 the binary variables in the first subset;
14 rounding up a first binary variable having a lowest ratio of the
15 cost penalty to the performance reward;
16 until no binary variables remain in a second subset, iteratively
17 performing the steps of:
18 determining the binary variables in the first subset that
19 may be rounded down without violating the performance
20 requirement, thereby forming the second subset;
21 calculating a cost reward and the performance reward
22 for each of the binary variables in the second subset;

23 rounding down one or more second binary variables in
24 the second subset having a zero performance reward;
25 rounding down a third binary variable in the second
26 subset corresponding to a highest ratio of a cost reward to
27 the performance reward if none of the binary variables in
28 the second subset have the zero performance reward; and
29 determining an upper bound for the minimum cost according to the
30 binary variables having binary values.

1 . 9. The method of claim 8 wherein the integer program further comprises a
2 . storage constraint.

1 10. The method of claim 9 wherein the step of determining the upper bound for
2 the minimum cost further comprises the steps of:

3 determining a particular node which uses a maximum amount of
4 storage within any evaluation interval; and
5 allocating the maximum amount of storage on all nodes for all
6 evaluation intervals.

1 11. The method of claim 9 wherein the step of determining the upper bound for
2 the minimum cost further comprises the steps of:

- 3 determining a maximum amount of storage for each node within any
- 4 evaluation interval; and
- 5 allocating the maximum amount of storage on each node for all
- 6 evaluation intervals.

1 12. The method of claim 8 wherein the integer program further comprises a
2 replica constraint.

1 13. The method of claim 12 wherein the step of determining the upper bound for
2 the minimum cost further comprises the steps of;
3 determining a maximum number of replicas for any data object within
4 any evaluation interval; and
5 placing the maximum number of replicas for all data objects for all

6 evaluation intervals.

1 14. The method of claim 12 wherein the step of determining the upper bound for
2 the minimum cost further comprises the steps of;

3 determining a maximum number of replicas for each data object within
4 any evaluation interval; and

5 placing the maximum number of replicas for each data object for all
6 evaluation intervals.

1 15. A computer readable memory comprising computer code for implementing a
2 method of determining bounds for a minimum cost, the method of determining the
3 bounds for the minimum cost comprising the steps of:

4 solving an integer program using a relaxation of binary variables to
5 determine a lower bound for the minimum cost, the integer program
6 comprising a performance constraint and an objective of minimizing a
7 cost, the binary variables having values between zero and one comprising
8 a first subset;

9 for the binary variables within the first subset and until no binary
10 variables remain in the first subset, iteratively performing the steps of:

11 rounding up a first binary variable having a lowest ratio of a
12 cost penalty to a performance reward; and

13 until no binary variables remain in a second subset, iteratively
14 performing the steps of:

15 determining the binary variables in the first subset that
16 may be rounded down without violating the performance
17 constraint, thereby forming the second subset;

18 rounding down one or more second binary variables in
19 the second subset having a zero performance reward; and

20 rounding down a third binary variable in the second
21 subset having a highest ratio of a cost reward to the
22 performance reward if none of the binary variables in the
23 second subset have the zero performance reward; and

24 determining an upper bound for the minimum cost according to the
25 binary variables having binary values.

1 16. The computer readable memory of claim 15 wherein the integer program
2 models a data placement problem.

1 17. The computer readable memory of claim 16 wherein the data placement
2 problem seeks to minimize the cost of placing data objects onto nodes of a
3 distributed storage system while meeting a performance requirement for a
4 workload.

1 18. The computer readable memory of claim 15 wherein the step of rounding up
2 the first binary variable within the subset further comprises calculating the cost
3 penalty and the performance reward.

1 19. The computer readable memory of claim 18 wherein the step of rounding
2 down the one or more second binary variables within the subset further comprises
3 calculating the performance reward.

1 20. The computer readable memory of claim 19 wherein the step of rounding
2 down the third binary variable within the subset further comprises calculating the
3 cost reward.

1 21. A computer readable memory comprising computer code for implementing a
2 method of determining bounds for a minimum cost, the method of determining the
3 bounds for the minimum cost comprising the steps of:

4 solving an integer program using a relaxation of binary variables to
5 determine a lower bound for the minimum cost, the relaxation allowing the
6 binary variables to take values over the range of zero to one, a first subset
7 of the binary variables comprising the binary variables having values
8 between zero and one, the integer program modeling a data placement
9 problem which seeks to minimize a cost of placing data objects onto nodes
10 of a distributed storage system while meeting a performance requirement
11 for a workload;

12 until no binary variables remain in the first subset, iteratively
13 performing the steps of:

14 calculating a cost penalty and a performance reward for each of
15 the binary variables in first the subset;
16 rounding up a first binary variable having a lowest ratio of the
17 cost penalty to the performance reward;
18 until no binary variables remain in a second subset, iteratively
19 performing the steps of:
20 determining the binary variables in the first subset that
21 may be rounded down without violating the performance
22 requirement, thereby forming the second subset;
23 calculating a cost reward and the performance reward
24 for each of the binary variables in the second subset;
25 rounding down one or more second binary variables in
26 the second subset having a zero performance reward;
27 rounding down a third binary variable in the second
28 subset corresponding to a highest ratio of a cost reward to
29 the performance reward if none of the binary variables in
30 the second subset have the zero performance reward; and
31 determining an upper bound for the minimum cost according to the
32 binary variables having binary values.

1 22. The computer readable memory of claim 21 wherein the integer program
2 further comprises a storage constraint.

1 23. The computer readable memory of claim 22 wherein the step of determining
2 the upper bound for the minimum cost further comprises the steps of:
3 determining a particular node which uses a maximum amount of
4 storage within any evaluation interval; and
5 allocating the maximum amount of storage on all nodes for all
6 evaluation intervals.

1 24. The computer readable memory of claim 22 wherein the step of determining
2 the upper bound for the minimum cost further comprises the steps of:
3 determining a maximum amount of storage for each node within any
4 evaluation interval; and

5 allocating the maximum amount of storage on each node for all
6 evaluation intervals.

1 25. The computer readable memory of claim 21 wherein the integer program
2 further comprises a replica constraint.

1 26. The computer readable memory of claim 25 wherein the step of determining
2 the upper bound for the minimum cost further comprises the steps of;
3 determining a maximum number of replicas for any data object within
4 any evaluation interval; and
5 placing the maximum number of replicas for all data objects for all
6 evaluation intervals.

1 27. The computer readable memory of claim 25 wherein the step of determining
2 the upper bound for the minimum cost further comprises the steps of;
3 determining a maximum number of replicas for each data object within
4 any evaluation interval; and
5 placing the maximum number of replicas for each data object for all
6 evaluation intervals.